arXiv:0905.3447v1 [math.OC] 21 May 2009

On convex problems in chance-constrained stochastic model predictive control ⋆

Eugenio Cinquemani [a], Mayank Agarwal [b], Debasish Chatterjee [a], John Lygeros [a]

[a] *Automatic Control Laboratory, Physikstrasse 3, ETH Zürich, 8092 Zürich, Switzerland*

[b] *Department of Electrical Engineering, Stanford University, California, USA*

Abstract

We investigate constrained optimal control problems for linear stochastic dynamical systems evolving in discrete time. We consider minimization of an expected value cost over a finite horizon. Hard constraints are introduced first, and then reformulated in terms of probabilistic constraints. It is shown that, for a suitable parametrization of the control policy, a wide class of the resulting optimization problems are convex, or admit reasonable convex approximations.

Key words: Stochastic control; Convex optimization; Probabilistic constraints

1 Introduction

This work stems from the attempt to address the optimal infinite-horizon constrained control of discrete-time stochastic processes by a model predictive control strategy [1,14,12,13,15,29,2,4,7]. We focus on linear dynamical systems driven by stochastic noise and a control input, and consider the problem of finding a control policy that minimizes an expected cost function while simultaneously fulfilling constraints on the control input and on the state evolution. In general, no control policy exists that guarantees satisfaction of deterministic (hard) constraints over the whole infinite horizon. One way to cope with this issue is to relax the constraints in terms of probabilistic (soft) constraints [25,26]. This amounts to requiring that constraints will not be violated with sufficiently large probability or, alternatively, that an expected reward for the fulfillment of the constraints is kept sufficiently large.

Two considerations lead to the reformulation of an infinite horizon problem in terms of subproblems of finite horizon length. First, given any bounded set (e.g. a safe set), the state of a linear stochastic dynamical system is guaranteed to exit the set at some time in the future with probability one whatever the control policy. Therefore, soft constraints may turn the original (infeasible) hard-constrained optimization problem into a feasible problem only if the horizon length is finite. Second, even if the constraints are reformulated so that an admissible infinite-horizon policy exists, the computation of such a policy is generally intractable. The aim of this note is to show that, for certain parameterizations of the policy space [21,6,17] and the constraints, the resulting finite horizon optimization problem is tractable.

An approach to infinite horizon constrained control problems that has proved successful in many applications is model predictive control [22]. In model predictive control, at every time t, a finite-horizon approximation of the infinite-horizon problem is solved but only the first control of the resulting policy is implemented. At the next time $t+1$, a measurement of the state is taken, a new finite-horizon problem is formulated, the control policy is updated,

⋆ This paper was not presented at any IFAC meeting. Research supported in part by the Swiss National Science Foundation under grant 200021- 122072. Corresponding author Eugenio Cinquemani. Tel. +41 (0)44 632 86 61; Fax +41 (0)44 632 12 11.

Email addresses: cinquemani@control.ee.ethz.ch (Eugenio Cinquemani), mayankag@stanford.edu (Mayank Agarwal), chatterjee@control.ee.ethz.ch (Debasish Chatterjee), lygeros@control.ee.ethz.ch (John Lygeros).

and the process is repeated in a receding horizon fashion. Under time-invariance assumptions, the finite-horizon optimal control problem is the same at all times, giving rise to a stationary optimal control policy that can be computed offline.

Motivated by the previous considerations, here we study the convexity of certain stochastic finite-horizon control problems with soft constraints. Convexity is central for the fast computation of the solution by way of numerical procedures, hence convex formulations [8] or convex approximations [23,9] of the stochastic control problems are commonly sought. However, for many of the classes of problems considered here, tight convex approximations are usually difficult to derive. One may argue that non-convex problems can be tackled by randomized algorithms [28,27,30]. However, randomized solutions are typically time-consuming and can only provide probabilistic guarantees. In particular, this is critical in the case where the system dynamics or the problem constraints are time-varying, since in that case optimization must be performed in real-time.

Here we provide conditions for the convexity of chance constrained stochastic optimal control problems. We derive and compare several explicit convex approximations of chance constraints for Gaussian noise processes and for polytopic and ellipsoidal constraint functions. Finally, we establish conditions for the convexity of a class of expectation-type constrains that includes standard integrated chance constraints [19,20] as a special case. For integrated chance constrains on Gaussian processes with polytopic constraint functions, an explicit formulation of the optimization problem is also derived.

The optimal constrained control problem we concentrate on is formulated in Section 2. A convenient parametrization of the control policies and the convexity of the objective function are discussed at this stage. Next, two probabilistic formulations of the constraints and conditions for the convexity of the space of admissible control policies are discussed: Section 3 is dedicated to chance constraints, while Section 4 is dedicated to integrated chance constraints. In Section 5, numerical simulations are reported to illustrate and discuss the results of the paper.

2 Problem statement

Let $\mathbb{N} = \{1, 2, \ldots\}$ and $\mathbb{N}_0 \triangleq \mathbb{N} \cup \{0\}$. Consider the following dynamical model: for $t \in \mathbb{N}_0$,

$$x(t+1) = Ax(t) + Bu(t) + w(t) \tag{1}$$

where $x(t) \in \mathbb{R}^n$ is the state, $u(t) \in \mathbb{R}^m$ is the control input, $A \in \mathbb{R}^{n\times n}$, $B \in \mathbb{R}^{n\times m}$, and $w(t)$ is a stochastic noise input defined on an underlying probability space $(\Omega, \mathfrak{F}, \mathbb{P})$. No assumption on the probability distribution of the process w is made at this stage. We assume that at any time $t \in \mathbb{N}_0$, $x(t)$ is observed exactly and that, for given $x_0 \in \mathbb{R}^n$, $x(0) = x_0$.

Fix a horizon length $N \in \mathbb{N}$. The evolution of the system from $t = 0$ through $t = N$ can be described in compact form as follows:

$$\bar{x} = \bar{A}x_0 + \bar{B}\bar{u} + \bar{D}\bar{w}, \tag{2}$$

where

$$\bar{x} \triangleq \begin{bmatrix} x(0) \\ x(1) \\ \vdots \\ x(N) \end{bmatrix}, \qquad \bar{u} \triangleq \begin{bmatrix} u(0) \\ u(1) \\ \vdots \\ u(N-1) \end{bmatrix}, \qquad \bar{w} \triangleq \begin{bmatrix} w(0) \\ w(1) \\ \vdots \\ w(N-1) \end{bmatrix}, \qquad \bar{A} \triangleq \begin{bmatrix} I_n \\ A \\ \vdots \\ A^N \end{bmatrix},$$

$$\bar{B} \triangleq \begin{bmatrix} 0_{n\times m} & \cdots & \cdots & 0_{n\times m} \\ B & \ddots & & \vdots \\ AB & B & \ddots & \vdots \\ \vdots & & \ddots & 0_{n\times m} \\ A^{N-1}B & \cdots & AB & B \end{bmatrix}, \qquad \bar{D} \triangleq \begin{bmatrix} 0_n & \cdots & \cdots & 0_n \\ I_n & \ddots & & \vdots \\ A & I_n & \ddots & \vdots \\ \vdots & & \ddots & 0_n \\ A^{N-1} & \cdots & A & I_n \end{bmatrix}.$$

Let $V : \mathbb{R}^{(N+1)n\times Nm} \to \mathbb{R}$ and $\eta : \mathbb{R}^{(N+1)n\times Nm} \to \mathbb{R}^r$, with $r \in \mathbb{N}$, be measurable functions. We are interested in constrained optimization problems of the following kind:

$$\begin{aligned} \inf_{\bar{u}\in\mathscr{U}} \quad & \mathbb{E}[V(\bar{x},\bar{u})] \\ \text{subject to} \quad & (2) \quad \text{and} \quad \eta(\bar{x},\bar{u}) \leq 0 \end{aligned} \tag{3}$$

where the expectation $\mathbb{E}[\cdot]$ is defined in terms of the underlying probability space $(\Omega, \mathfrak{F}, \mathbb{P})$, $\mathscr{U}$ is a class of causal deterministic state-feedback control policies and the inequality in (3) is interpreted componentwise.

Example 1 In the (unconstrained) linear stochastic control problem [3], w is Gaussian white noise and the aim is to minimize

$$\mathbb{E}\left[\sum_{t=0}^{N-1}\big(x^T(t)Q(t)x(t) + u^T(t)R(t)u(t)\big) + x^T(N)Q(N)x(N)\right],$$

where the matrices $Q(t) \in \mathbb{R}^{n\times n}$ and $R(t) \in \mathbb{R}^{m\times m}$ are positive definite for all t, with respect to causal feedback policies subject to the system dynamics (1). This problem fits easily in our framework; it suffices to define

$$V(\bar{x},\bar{u}) = \begin{bmatrix}\bar{x}^T & \bar{u}^T\end{bmatrix} M \begin{bmatrix}\bar{x}\\ \bar{u}\end{bmatrix}, \tag{4}$$

with $M = \operatorname{diag}\big(Q(0), Q(1), \ldots, Q(N), R(0), R(1), \ldots, R(N-1)\big) > 0$ (the notation $M > 0$ indicates that M is a positive definite matrix). In our framework, though, the input noise sequence may have an arbitrary correlation structure, the cost function may be non-quadratic and, most importantly, constraints may be present.

Standard constraints on the state and the input are also formulated easily. For instance, sequential ellipsoidal constraints of the type

$$\begin{aligned} \begin{bmatrix}x^T(t) & u^T(t)\end{bmatrix} S(t) \begin{bmatrix}x(t)\\ u(t)\end{bmatrix} \leq 1, \qquad & t = 0, 1, \ldots, N-1, \\ x^T(N)S(N)x(N) \leq 1, & \end{aligned}$$

with $0 < S(t) \in \mathbb{R}^{(n+m)\times(n+m)}$ for $t = 0, 1, \ldots, N-1$ and $0 < S(N) \in \mathbb{R}^{n\times n}$, are captured by the definition

$$\eta(\bar{x},\bar{u}) = \begin{bmatrix}\eta_0(\bar{x},\bar{u}) & \eta_1(\bar{x},\bar{u}) & \cdots & \eta_N(\bar{x})\end{bmatrix}^T$$

where, for $t = 0, 1, \ldots, N$,

$$\eta_t(\bar{x},\bar{u}) = \begin{bmatrix}\bar{x}^T & \bar{u}^T\end{bmatrix} \Xi_t \begin{bmatrix}\bar{x}\\ \bar{u}\end{bmatrix} - 1,$$

and each matrix Ξ_t is immediately constructed in terms of the $S(t)$. Our framework additionally allows for cross-constraints between states and inputs at different times.

2.1 Feedback from the noise input

By the hypothesis that the state is observed without error, one may reconstruct the noise sequence from the sequence of observed states and inputs by the formula

$$w(t) = x(t+1) - Ax(t) - Bu(t), \qquad t \in \mathbb{N}_0. \tag{5}$$

In light of this, and following [21,17,6], we shall consider policies of the form:

$$u(t) = \sum_{i=0}^{t-1} G_{t,i}w(i) + d_t, \tag{6}$$

where the feedback gains $G_{t,i} \in \mathbb{R}^{m\times n}$ and the affine terms $d_t \in \mathbb{R}^m$ must be chosen based on the control objective. With this definition, the value of u at time t depends on the values of w up to time $t-1$. Using (5) we see that $u(t)$ is a function of the observed states up to time t. It was shown in [17] that there exists a (nonlinear) bijection between control policies in the form (6) and the class of affine state feedback policies. That is, provided one is interested in affine state feedback policies, parametrization (5) constitutes no loss of generality. Of course, this choice is generally suboptimal, since there is no reason to expect that the optimal policy is affine, but it will ensure the tractability of a large class of optimal control problems. In compact notation, the control sequence up to time $N-1$ is given by

$$\bar{u} = \bar{G}\bar{w} + \bar{d}, \tag{7}$$

where

$$\bar{G} \triangleq \begin{bmatrix} 0_{m\times n} & & & \\ G_{1,0} & 0_{m\times n} & & \\ \vdots & \ddots & \ddots & \\ G_{N-1,0} & \cdots & G_{N-1,N-2} & 0_{m\times n} \end{bmatrix}, \qquad \bar{d} \triangleq \begin{bmatrix} d_0 \\ d_1 \\ \vdots \\ d_{N-1} \end{bmatrix}; \tag{8}$$

note the lower triangular structure of $\bar{G}$ that enforces causality. The resulting closed-loop system dynamics can be written compactly as the equality constraint

$$\bar{x} = \bar{A}x_0 + \bar{B}(\bar{G}\bar{w} + \bar{d}) + \bar{D}\bar{w}. \tag{9}$$

Let us denote the parameters of the control policy by $\theta = (\bar{G}, \bar{d})$ and write $(\bar{x}_\theta, \bar{u}_\theta)$ to emphasize the dependence of $\bar{x}$ and $\bar{u}$ on θ. From now on we will consider the optimization problem

$$\begin{aligned} \inf_{\theta\in\Theta} \quad & \mathbb{E}[V(\bar{x}_\theta, \bar{u}_\theta)] && (10)\\ \text{subject to} \quad & (7),\ (9) \text{ and} && (11)\\ & \eta(\bar{x}_\theta, \bar{u}_\theta) \le 0, && (12) \end{aligned}$$

where Θ is the linear space of optimization parameters in the form (8).

Remark 2 With the above parametrization of the control policy, both $\bar{u}_\theta$ and $\bar{x}_\theta$ are affine functions of the parameters θ (for fixed $\bar{w}$) and of the process noise $\bar{w}$ (for fixed θ). Most of the results developed below rely essentially on this property. It was noticed in [5] that a parametric causal feedback control policy with the same property can be easily defined based on indirect observations of the state, provided the measurement model is linear. The method enables one to extend the results of this paper to the case of linear output feedback. For the sake of conciseness, this extension will not be pursued here.

2.2 Optimal control problem with relaxed constraints

In general, no control policy can ensure that the constraint (12) is satisfied for all outcomes of the stochastic input $\bar{w}$. In the standard LQG setting, for instance, any nontrivial constraint on the system state would be violated with nonzero probability. We therefore consider relaxed formulations of the constrained optimization problem (10)–(12) of the form

$$\begin{aligned} \inf_{\theta\in\Theta} \quad & \mathbb{E}[V(\bar{x}_\theta, \bar{u}_\theta)] && (13)\\ \text{subject to} \quad & (7),\ (9) \text{ and} && (14)\\ & \mathbb{E}[\phi \circ \eta(\bar{x}_\theta, \bar{u}_\theta)] \le 0, && (15) \end{aligned}$$

where $\phi : \mathbb{R}^r \to \mathbb{R}^R$, with $R \in \mathbb{N}$, is a convenient measurable function and the inequality is again interpreted componentwise. For appropriate choices of ϕ, this formulation embraces most common probabilistic constraint relaxations, including chance constraints (see e.g. [23]), integrated chance constraints [19,20], and expectation constraints (see e.g. [26]).

We are interested in the convexity of the optimization problem (13)–(15). First we establish a general convexity result.

Proposition 3 *Let $(\Omega, \mathfrak{F}, \mathbb{P})$ be a probability space, Θ be a convex subset of a vector space and $\mathscr{D} \subseteq \mathbb{R}$ be convex. Let $\gamma : \Omega \times \Theta \to \mathscr{D}$ and $\varphi : \mathscr{D} \to \mathbb{R}$ be measurable functions and define*

$$J(\theta) \triangleq \mathbb{E}[\varphi \circ \gamma(\omega, \theta)].$$

Assume that:

(i) *the mapping $\gamma(\omega, \cdot) : \Theta \to \mathbb{R}$ is convex for almost all $\omega \in \Omega$;*
(ii) *φ is monotone nondecreasing and convex;*
(iii) *$J(\theta)$ is finite for all $\theta \in \Theta$.*

Then the mapping $J : \Theta \to \mathbb{R}$ is convex.

Proof: Fix a generic $\omega \in \Omega$. Since $\gamma(\omega, \theta)$ is convex in θ and φ is monotone nondecreasing, for any $\theta, \theta' \in \Theta$ and any $\alpha \in [0, 1]$,

$$\varphi\big(\gamma(\omega, \alpha\theta + (1-\alpha)\theta')\big) \leq \varphi\big(\alpha\gamma(\omega, \theta) + (1-\alpha)\gamma(\omega, \theta')\big).$$

Moreover, since φ is convex,

$$\varphi\big(\alpha\gamma(\omega, \theta) + (1-\alpha)\gamma(\omega, \theta')\big) \leq \alpha\varphi\big(\gamma(\omega, \theta)\big) + (1-\alpha)\varphi\big(\gamma(\omega, \theta')\big).$$

Since these inequalities hold for almost all $\omega \in \Omega$, it follows that

$$\begin{aligned}\mathbb{E}[\varphi\big(\gamma(\omega, \alpha\theta + (1-\alpha)\theta')\big)] &\leq \mathbb{E}[\alpha\varphi\big(\gamma(\omega, \theta)\big) + (1-\alpha)\varphi\big(\gamma(\omega, \theta')\big)] \\ &= \alpha\mathbb{E}[\varphi\big(\gamma(\omega, \theta)\big)] + (1-\alpha)\mathbb{E}[\varphi\big(\gamma(\omega, \theta')\big)],\end{aligned}$$

which proves the assertion. □

Assumption (iii) can be replaced by either of the following:

(iii′) $J(\theta) \in \mathbb{R} \cup \{+\infty\}$, $\forall \theta \in \Theta$.
(iii″) $J(\theta) \in \mathbb{R} \cup \{-\infty\}$, $\forall \theta \in \Theta$.

Let us now make the following standing assumption.

Assumption 1 *$V(\bar{x}, \bar{u})$ is a convex function of $(\bar{x}, \bar{u})$ and $\mathbb{E}[V(\bar{x}_\theta, \bar{u}_\theta)]$ is finite for all $\theta \in \Theta$.*

Proposition 4 *Under Assumption 1, $\mathbb{E}[V(\bar{x}_\theta, \bar{u}_\theta)]$ is a convex function of θ.*

Proof: First, note that the set Θ of admissible parameters θ is a linear space. Let us write $\bar{w}(\omega)$ $\bar{x}_\theta(\omega)$ and $\bar{u}_\theta(\omega)$ to express the dependence of $\bar{w}$, $\bar{x}_\theta$ and $\bar{u}_\theta$ on the random event $\omega \in \Omega$. Fix ω arbitrarily. Since the mapping

$$\theta \mapsto \begin{bmatrix}\bar{x}_\theta(\omega) \\ \bar{u}_\theta(\omega)\end{bmatrix} = \begin{bmatrix}\bar{A}x_0 + \bar{B}(\bar{G}\bar{w}(\omega) + \bar{d}) + \bar{D}\bar{w}(\omega) \\ \bar{G}\bar{w}(\omega) + \bar{d}\end{bmatrix}$$

is affine and the mapping $(\bar{x}, \bar{u}) \mapsto V(\bar{x}, \bar{u})$ is assumed convex, their combination $\theta \mapsto V(\bar{x}_\theta, \bar{u}_\theta)$ is a convex function of θ. Then, the result follows from Proposition 3 with $\gamma(\omega, \theta) = V\big(\bar{x}_\theta(\omega), \bar{u}_\theta(\omega)\big)$ and φ equal to the identity map. □

By virtue of the alternative assumptions (iii′) and (iii″) of Proposition 3, the requirement that $\mathbb{E}[V(\bar{x}_\theta, \bar{u}_\theta)]$ be finite for all θ may be relaxed. A sufficient requirement is that there exist no two values θ and θ' such that $J_{x_0}(\theta) = +\infty$ and $J_{x_0}(\theta') = -\infty$. In particular, the result applies to quadratic cost functions of the type (4), with $M \geq 0$.

In general, the relaxed constraint (15) is nonconvex even if the components $\eta_i : \mathbb{R}^{(N+1)n \times Nm} \to \mathbb{R}$, with $i = 1, \ldots, r$, of the vector function η are convex. In the next sections we will study the convexity and provide convex approximations of (15) for different approaches to probabilistic relaxation of hard constraints, i.e. for different choices of the function ϕ.

3 Chance Constraints

For a given $\alpha \in]0,1[$, we relax the hard constraint $\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0$ by requiring that it be satisfied with probability $1-\alpha$. Hence we address the optimization problem

$$\begin{aligned} \inf_{\theta\in\Theta} \quad & \mathbb{E}[V(\bar{x}_\theta, \bar{u}_\theta)] && (16)\\ \text{subject to} \quad & (7),\ (9) \text{ and} && (17)\\ & \mathbb{P}(\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0) \geq 1-\alpha. && (18) \end{aligned}$$

The smaller α, the better the approximation of the hard constraint (12) at the expense of a more constrained optimization problem. This problem is obtained as a special case of Problem (13)–(15) by setting $R=1$ and defining ϕ as

$$\phi_{CC}(\eta) = 1 - \prod_{i=1}^{r} \mathbf{1}_{]-\infty,0]}(\eta_i) - \alpha,$$

where $\mathbf{1}_{]-\infty,0]}(\cdot)$ is the standard indicator function. We now study the convexity of (18) with respect to θ.

3.1 The fixed feedback case

First assume that the feedback term $\bar{G}$ in (7) is fixed and consider the convexity of the optimization problem (16)–(18) with respect to the open loop control action $\bar{d}$. That is, for a given $\bar{G}$ in the form (8), the parameter space Θ becomes the set $\{(\bar{G}, \bar{d}),\ \forall \bar{d} \in \mathbb{R}^{Nm}\}$. For the given $\bar{G}$ and $i = 1, \ldots, r$ define

$$g_i(\bar{d}, \bar{w}) = \eta_i(\bar{A}x_0 + \bar{B}(\bar{G}\bar{w} + \bar{d}) + \bar{D}\bar{w}, \bar{G}\bar{w} + \bar{d}),$$

where $\eta_i : \mathbb{R}^{(N+1)n+Nm} \to \mathbb{R}$ is the i-th element of the constraint function η. Define

$$p(\bar{d}) = \mathbb{P}[g_1(\bar{d}, \bar{w}) \leq 0, \ldots g_r(\bar{d}, \bar{w}) \leq 0] \tag{19}$$

and $\mathscr{F}_{CC} \triangleq \{\bar{d} :\ p(\bar{d}) \geq 1-\alpha\}$. Observe that $\mathscr{F}_{CC}$ corresponds to the constraint set dictated by (17)–(18) when $\bar{G}$ is fixed.

Proposition 5 *Assume that $\bar{w}$ has a continuous distribution with log-concave probability density and that, for $i = 1, \ldots, r$, $g_i : \mathbb{R}^{Nm+Nn} \to \mathbb{R}$ is quasi-convex. Then, for any value of $\alpha \in]0,1[$, $\mathscr{F}_{CC}$ is convex. As a consequence, under Assumption 1 and for any $\alpha \in]0,1[$, the optimization problem*

$$\begin{aligned} \inf_{\theta\in\Theta} \quad & \mathbb{E}[V(\bar{x}_\theta, \bar{u}_\theta)]\\ \text{subject to} \quad & (7),\ (9) \text{ and } p(\bar{d}) \geq 1-\alpha \end{aligned}$$

is convex.

Proof: It follows from [24, Theorem 10.2.1] that (19) is a log-concave function of $\bar{d}$, i.e. the mapping $\bar{d} \mapsto \log p(\bar{d})$ is concave. Since log is a monotone increasing function, we may write that $\mathscr{F}_{CC} = \{\bar{d} :\ \log p(\bar{d}) \geq \log(1-\alpha)\}$. Hence, $\mathscr{F}_{CC}$ is a convex set. The convexity of the optimization problem follows readily from Assumption 1. □

Among others, Gaussian, exponential and uniform distributions are continuous with log-concave probability density. As for the functions g_i, one case of interest where the assumptions of Proposition 5 are fulfilled is when $\eta(\bar{x}, \bar{u})$ is affine in $\bar{x}$ and $\bar{u}$. This is the case of polytopic constraints, which will be treated extensively in the next section. Apparently, this convexity result cannot be applied to the ellipsoidal constraints treated subsequently in Section 3.3, nor can it be extended to the general constraint (18) with both $\bar{d}$ and $\bar{G}$ varying. Loosely speaking, the latter is because the functions g_i are not simultaneously quasi-concave in $\bar{w}$ and $\bar{G}$. In the next sections we will develop convex conservative approximations of (18) for various definitions of η.

3.2 Polytopic Constraint Functions

Throughout the rest of Section 3 we shall rely on the following assumption.

Assumption 2 *$\bar{w}$ is a Gaussian random vector with mean zero and covariance matrix $\bar{\Sigma} > 0$, denoted by $\bar{w} \sim \mathcal{N}(0, \bar{\Sigma})$.*

Polytopic constraint functions

$$\eta(\bar{x}_\theta, \bar{u}_\theta) = T^x \bar{x}_\theta + T^u \bar{u}_\theta - y, \tag{20}$$

where $T^x \in \mathbb{R}^{r \times (N+1)n}, T^u \in \mathbb{R}^{r \times Nm}$, and $y \in \mathbb{R}^r$, describe one of the most common types of constraints. In light of (7) and (9),

$$\eta(\bar{x}_\theta, \bar{u}_\theta) = h_\theta + P_\theta \bar{w}, \tag{21}$$

where $h_\theta = (T^x \bar{A} x_0 - y) + (T^x \bar{B} + T^u)\bar{d}$, and $P_\theta = (T^x \bar{D} + (T^x \bar{B} + T^u)\bar{G})$. It is thus apparent that $\eta(\bar{x}_\theta, \bar{u}_\theta)$ is affine in the parameters θ. Yet, in general, constraint (18) is nonconvex. We now describe three approaches to approximate constraints (18) that lead to convex conservative constraints.

3.2.1 Approximation via constraint separation

Constraint (18) requires us to satisfy $\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0$ with probability of at least $1 - \alpha$. Here $\eta \in \mathbb{R}^r$ and the inequality $\eta \leq 0$ is interpreted componentwise. One idea is to select coefficients $\alpha_i \in (0, 1)$ such that $\sum_{i=1}^r \alpha_i = \alpha$ and to satisfy the inequalities $\mathbb{P}(\eta_i \leq 0) \geq 1 - \alpha_i$, with $i = 1, \ldots, r$. Note that this choice is obtained in (15) by setting $R = r$ and $\phi_i(\eta) = 1 - \mathbf{1}_{]-\infty, 0]}(\eta_i) - \alpha_i$ where, for $i = 1, \ldots, r$, $\phi_i : \mathbb{R}^r \to \mathbb{R}$ denotes the i-th component of function ϕ.

Let $h_{i,\theta}$ and $P_{i,\theta}^T$ be the i-th entry of h_θ and the i-th row of P_θ, respectively, and let $\bar{\Sigma}^{\frac{1}{2}}$ be a symmetric real matrix square root of $\bar{\Sigma}$.

Proposition 6 *Let $\alpha_i \in]0, 1[$. Under Assumption 2, the constraint*

$$\mathbb{P}(\eta_i(\bar{x}_\theta, \bar{u}_\theta) \leq 0) \geq 1 - \alpha_i,$$

with η defined as in (20), is equivalent to the second-order cone constraint in the parameters $\theta \in \Theta$

$$h_{i,\theta} + \beta_i \left\| \bar{\Sigma}^{\frac{1}{2}} P_{i,\theta} \right\| \leq 0$$

where $\beta_i = \sqrt{2}\,\mathrm{erf}^{-1}(1 - 2\alpha_i)$ and $\mathrm{erf}^{-1}(\cdot)$ is the inverse of the standard error function $\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x \mathrm{e}^{-u^2}\,\mathrm{d}u$. As a consequence, under Assumption 1 and if $\alpha_1 + \ldots + \alpha_r = \alpha$, the problem

$$\begin{aligned} \inf_{\theta \in \Theta} \quad & \mathbb{E}[V(\bar{x}_\theta, \bar{u}_\theta)] \\ \text{subject to} \quad & (7),\ (9) \text{ and } h_{i,\theta} + \beta_i \left\| \bar{\Sigma}^{\frac{1}{2}} P_{i,\theta} \right\| \leq 0,\ i = 1, \ldots, r \end{aligned}$$

is a convex conservative approximation of Problem (16)–(18).

Proof: From Eq. (21) we can write $\eta_i = h_{i,\theta} + P_{i,\theta}^T \bar{w}$. Since $\bar{w} \sim \mathcal{N}(0, \bar{\Sigma})$, η_i is also Gaussian with distribution $\mathcal{N}(h_{i,\theta}, P_{i,\theta}^T \bar{\Sigma} P_{i,\theta})$. It is easily seen that, for any scalar Gaussian random variable X with distribution $\mathcal{N}(\mu, \sigma^2)$,

$$\mathbb{P}(X \leq 0) \geq 1 - \alpha \iff \mu + \beta\sigma \leq 0,$$

where $\beta = \sqrt{2}\,\mathrm{erf}^{-1}(1 - 2\alpha)$. Hence the constraint $\mathbb{P}(\eta_i(\bar{x}_\theta, \bar{u}_\theta) \leq 0) \geq 1 - \alpha_i$ is equivalent to $h_{i,\theta} + \beta_i \left\| \bar{\Sigma}^{\frac{1}{2}} P_{i,\theta} \right\| \leq 0$, where $\beta_i = \sqrt{2}\,\mathrm{erf}^{-1}(1 - 2\alpha_i)$. Since $h_{i,\theta}$ and $P_{i,\theta}$ are both affine in the parameters $\theta = (\bar{G}, \bar{d})$, the above constraint is a second-order cone constraint in θ. It is easy to see that this choice guarantees that $\mathbb{P}(\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0) \geq 1 - \alpha$. □

3.2.2 Approximation via confidence ellipsoids

The approach of Section 3.2.1 may be too conservative since the probability of a union of events is approximated by the sum of the probabilities of the individual events. One can also calculate a conservative approximation of the union at once. Constraint (18) with $\eta(\bar{x}_\theta, \bar{u}_\theta)$ as in (21) restricts the choice of P_θ and h_θ to be such that, with a probability of $1-\alpha$ or more, the realization of random vector $\eta(\bar{x}_\theta, \bar{u}_\theta)$ lies in the negative orthant $\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0$. In general, it is difficult to describe this constraint explicitly since it involves the integration of the probability density of η over the negative orthant. However, an explicit approximation of the constraint can be computed by ensuring that the $100(1-\alpha)\%$ confidence ellipsoid of η is contained in the negative orthant. Fulfilling this requirement automatically implies that the probability of $\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0$ is strictly greater than $1-\alpha$.

Since $\bar{w} \sim \mathcal{N}(0, \bar{\Sigma})$, it follows that $\eta(\bar{x}_\theta, \bar{u}_\theta) = h_\theta + P_\theta \bar{w} \sim \mathcal{N}(h_\theta, \bar{\Sigma}_\theta)$, with $\bar{\Sigma}_\theta := P_\theta \bar{\Sigma} P_\theta^T$. Consider the case where $\bar{\Sigma}_\theta$ is invertible. Define the r-dimensional ellipsoid

$$\mathcal{E}(h_\theta, \bar{\Sigma}_\theta, \beta) = \big\{\eta \in \mathbb{R}^r : \ | \ (\eta - h_\theta)^T \bar{\Sigma}_\theta^{-1} (\eta - h_\theta) \leq \beta^2\big\}, \tag{22}$$

where $\beta > 0$ is a parameter specifying the size of the ellipsoid. Notice that, in general, $\bar{\Sigma}_\theta$ is invertible when $r \leq Nn$ (i.e. the number of constraints is less than the total dimension of the process noise). If $r > Nn$, as there are Nn independent random variables in the optimization problem, the following result still holds with Nn in place of r.

Proposition 7 *Let $\alpha \in]0,1[$. Under Assumption 2, the constraint $\mathbb{P}[\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0] \geq 1-\alpha$ with η defined as in* (20) *is conservatively approximated by the constraint*

$$\mathcal{E}\big(h_\theta, \bar{\Sigma}_\theta, \beta(\alpha)\big) \subset (-\infty, 0\,]^{\,r}, \tag{23}$$

where $\beta(\alpha) = \sqrt{F^{-1}(1-\alpha)}$ and $F(\cdot)$ is the probability distribution function of a χ^2 random variable with r degrees of freedom. Moreover, (23) *can be reformulated as the set of second-order cone constraints*

$$h_{i,\theta} + \beta(\alpha) \left\| \bar{\Sigma}^{\frac{1}{2}} P_{i,\theta} \right\| \leq 0, \quad i = 1, \ldots, r. \tag{24}$$

As a consequence, under Assumption 1, the problem

$$\begin{aligned} \inf_{\theta \in \Theta} \quad & \mathbb{E}[V(\bar{x}_\theta, \bar{u}_\theta)] \\ \text{subject to} \quad & (7),\ (9) \text{ and } (24) \end{aligned}$$

is a convex conservative approximation of Problem (16)–(18).

Proof: Since $\eta(\bar{x}_\theta, \bar{u}_\theta) \sim \mathcal{N}(h_\theta, \bar{\Sigma}_\theta)$, the random variable

$$\big(\eta(\bar{x}_\theta, \bar{u}_\theta) - h_\theta\big)^T \bar{\Sigma}_\theta^{-1} \big(\eta(\bar{x}_\theta, \bar{u}_\theta) - h_\theta\big)$$

is χ^2 with r degrees of freedom. Then, choosing β such that $F(\beta^2) = 1-\alpha$ guarantees that $\mathcal{E}\big(h_\theta, \bar{\Sigma}_\theta, \beta(\alpha)\big)$ is the $100(1-\alpha)\%$ confidence ellipsoid for $\eta(\bar{x}_\theta, \bar{u}_\theta)$. Finally, under (23), $\mathbb{P}[\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0] \geq \mathbb{P}[\eta(\bar{x}_\theta, \bar{u}_\theta) \in \mathcal{E}\big(h_\theta, \bar{\Sigma}_\theta, \beta(\alpha)\big)] = 1-\alpha$, which proves the first claim. To prove the second claim, note that (22) can alternatively be represented as

$$\mathcal{E}\big(h_\theta, \bar{\Sigma}_\theta, \beta\big) = \big\{\eta \in \mathbb{R}^r \ \big| \ \eta = h_\theta + M_\theta u,\ \|u\| \leq 1\big\}, \tag{25}$$

where $M_\theta = \beta \bar{\Sigma}_\theta^{\frac{1}{2}}$. Since $\eta \leq 0$ if and only if $e_i^T \eta \leq 0$, $i = 1, \cdots, r$, where the e_i denote the standard basis vectors in $\mathbb{R}^r$, we may rewrite (23) as

$$e_i^T (h_\theta + M_\theta u) \leq 0 \quad \forall \, \|u\| \leq 1,$$

or equivalently

$$\sup_{\|u\| \leq 1} e_i^T (h_\theta + M_\theta u) \leq 0$$

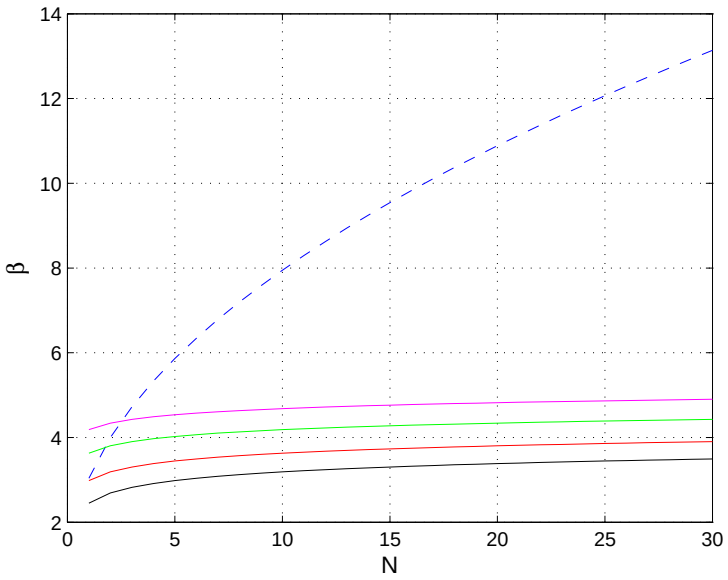

Fig. 1. Values of β as a function of the horizon length N for the ellipsoidal method (dashed line) and the constraint separation method (solid lines). It is assumed that $r = N \cdot \bar{r} \cdot (n+m)$, where $n = 5$ is the dimension of the system state and of the process noise, $m = 2$ is the dimension of the input u. $\bar{r}$ reflects the number of constraints per stage; we take $\bar{r} = 2, 10, 100, 1000$. The value of β is invariant with respect to $\bar{r}$ (single dashed line) for the ellipsoidal method, while it increases with $\bar{r}$ (multiple solid lines) for the constraint separation method.

for $i = 1, \cdots, r$. For each i, the supremum is attained with $u = M_\theta^T e_i / \left\| M_\theta^T e_i \right\|$; therefore, the above is equivalent to $e_i^T h_\theta + \left\| M_\theta^T e_i \right\| \leq 0$. Clearly, $e_i^T h_\theta = h_{i,\theta}$. Moreover, since $M_\theta = \beta \bar{\Sigma}_\theta^{\frac{1}{2}}$, we have

$$\left\| M_\theta^T e_i \right\| = \beta \sqrt{e_i^T \bar{\Sigma}_\theta e_i} = \beta \sqrt{e_i^T P_\theta \bar{\Sigma} P_\theta^T e_i} = \beta \left\| \bar{\Sigma}^{\frac{1}{2}} P_{i,\theta} \right\|.$$

Therefore, constraint (23) reduces to (24). Since the variables $h_{i,\theta}$ and $P_{i,\theta}$ are affine in the original parameters $\theta = (\bar{G}, \bar{d})$, this is an intersection of second order cone constraints. As a result, under the additional Assumption 1, the optimization of $\mathbb{E}[V(\bar{x}_\theta, \bar{u}_\theta)]$ with (24) in place of (18) is a convex conservative approximation of (16)–(18). □

3.2.3 Comparison of constraint separation and confidence ellipsoid methods

In light of Propositions 6 and 7, both approaches lead to formulating constraints of the form

$$\mu_i + \beta \sigma_i \leq 0,$$

with $i = 1, \ldots, r$, where μ_i and σ_i are the mean and the standard deviation of the scalar random variables $\eta_i(\bar{x}_\theta, \bar{u}_\theta)$, and smaller values of the constant β correspond to less conservative approximations of the original chance constraint. For a given value of α, the value of β depends on the number of constraints r in a way that differs in the two cases. In the confidence ellipsoid method, in particular, the value of β is determined by Nn (the total dimension of the process noise) when $r \geq Nn$. In Figure 1, we compare the growth of β in the two approaches under the assumption that r grows linearly with the horizon length N. (For the constraint separation method we choose $\alpha_1 = \ldots = \alpha_r = \alpha/r$, so that the value of β is the same for all constraints.) The increase of β is quite rapid in the confidence ellipsoid method, which is only effective for a small number of constraints. An explanation of this phenomenon is provided by the following fact, that is better known by the name of (classical) "concentration of measure" inequalities; proofs may be found in, e.g., [10].

Proposition 8 *Let $\Gamma_{h,\Sigma'}$ be the r-dimensional Gaussian measure with mean h and (nonsingular) covariance Σ', i.e.,*

$$\Gamma_{h,\Sigma'}(\mathrm{d}\xi) = \frac{1}{(2\pi)^{r/2}\sqrt{\det \Sigma'}} \exp\left(-\frac{1}{2}\left\langle \xi - h, \Sigma'^{-1}(\xi - h)\right\rangle\right)\mathrm{d}\xi.$$

Then for $\varepsilon \in]0,1[$,

(i) $\Gamma_{h,\Sigma'}\left(\left\{\xi \;\middle|\; \|\xi - h\|_{\Sigma'^{-1}} > \sqrt{\frac{r}{1-\varepsilon}}\right\}\right) \leq \mathrm{e}^{-\frac{r\varepsilon^2}{4}}$;

(ii) $\Gamma_{h,\Sigma'}\left(\left\{\xi \middle|\, \|\xi - h\|_{\Sigma'^{-1}} < \sqrt{r(1-\varepsilon)}\right\}\right) \leq \mathrm{e}^{-\frac{r\varepsilon^2}{4}}$.

The above proposition states that as the dimension r of the Gaussian measure increases, its mass concentrates in an ellipsoidal shell of 'mean-size' $\sqrt{r}$. It readily follows that since $\eta(\bar{x}_\theta, \bar{u}_\theta)$ is a r-dimensional Gaussian random vector, its mass concentrates around a shell of size $\sqrt{r}$. Note that the bounds corresponding to (i) and (ii) of Proposition 8 in the case of $\eta(\bar{x}_\theta, \bar{u}_\theta)$ are independent of the optimization parameters θ; of course the relative sizes of the confidence ellipsoids change with θ (because the mean and the covariance of $\eta(\bar{x}_\theta, \bar{u}_\theta)$ depend on θ), but Proposition 8 shows that the size of the confidence ellipsoids grow quite rapidly with the dimension of the noise and the length of the optimization horizon. Intuitively one would expect the ellipsoidal constraint approximation method to be more effective than the cruder approximation by constraint separation. Figure 1 and Proposition 8 however suggest that this is not the case in general; for large numbers of constraints (e.g. longer MPC prediction horizon) the constraint separation method is the less conservative.

3.2.4 Approximation via expectations

For any r-dimensional random vector η, we have $1 - \mathbb{P}\big(\eta \leq 0\big) = \mathbb{E}\left[1 - \prod_{i=1}^r \mathbf{1}_{]-\infty,0]}(\eta_i)\right]$. Using this fact one can arrive at conservative convex approximations of the chance-constraint (18) by replacing the function in the expectation with appropriate approximating functions. For $t_i > 0$, $i = 1, \ldots, r$, consider

$$\varphi(\eta) = \sum_{i=1}^r \exp(t_i \eta_i).$$

Lemma 9 *For any r-dimensional random vector η, $\mathbb{E}[\varphi(\eta)] \geq 1 - \mathbb{P}[\eta \leq 0]$.*

Proof: For every fixed value of η, it holds that $\varphi(\eta) \geq 1 - \prod_{i=1}^r \mathbf{1}_{]-\infty,0]}(\eta_i)$. Hence $\mathbb{E}[\varphi(\eta)] \geq \mathbb{E}[1 - \prod_{i=1}^r \mathbf{1}_{]-\infty,0]}(\eta_i)] = 1 - \mathbb{P}[\eta \leq 0]$. □

Proposition 10 *Under Assumption 2, for η defined as in* (20)*, it holds that*

$$\mathbb{E}\big[\varphi\big(\eta(\bar{x}_\theta, \bar{u}_\theta)\big)\big] = \sum_{i=1}^r \exp\left(t_i h_{i,\theta} + \frac{t_i^2}{2}||\bar{\Sigma}^{\frac{1}{2}} P_{i,\theta}||^2\right).$$

As a consequence, under Assumption 1 and for any choice of $t_i > 0$, $i = 1, \ldots, r$, the problem

$$\begin{aligned} \inf_{\theta \in \Theta} \quad & \mathbb{E}[V(\bar{x}_\theta, \bar{u}_\theta)] \\ \text{subject to} \quad & (7),\ (9) \text{ and } \sum_{i=1}^r \exp\left(t_i h_{i,\theta} + \frac{t_i^2}{2}||\bar{\Sigma}^{\frac{1}{2}} P_{i,\theta}||^2\right) \leq \alpha \end{aligned}$$

is a convex conservative approximation of Problem (16)–(18)*.*

Proof: It is easily seen that, for any r-dimensional Gaussian random vector η with mean μ and covariance matrix Σ', and any vector $c \in \mathbb{R}^r$, $\mathbb{E}\big[\exp(c^T\eta)\big] = \exp\big(c^T\mu + \frac{1}{2}c^T\Sigma' c\big)$. Let us now write η_θ in place of $\eta(\bar{x}_\theta, \bar{u}_\theta)$ for shortness. By the hypotheses on $\bar{w}$, in the light of (21), η_θ is Gaussian with mean h_θ and covariance $P_\theta \bar{\Sigma} P_\theta^T$. Then, for a vector $c_i \in \mathbb{R}^r$ with zero entries except for a coefficient t_i in the i-th position,

$$\mathbb{E}[\exp(c_i^T\eta_\theta)] = \exp\Big(c_i^T h_\theta + \frac{1}{2}c_i^T P_\theta \bar{\Sigma} P_\theta^T c_i\Big) = \exp\Big(t_i h_{i,\theta} + \frac{t_i^2}{2} P_{i,\theta}^T \bar{\Sigma} P_{i,\theta}\Big).$$

Summing up for $i = 1, \dots, r$ yields the first result. In order to prove the second statement, note that

$$\begin{aligned}\mathbb{E}\big[\varphi(\eta_\theta)\big] \le \alpha &\iff \log \mathbb{E}\big[\varphi(\eta_\theta)\big] \le \log\alpha \\ &\iff \log\left(\sum_{i=1}^r \exp\Big(t_i h_{i,\theta} + \frac{t_i^2}{2}\left\|\bar{\Sigma}^{\frac{1}{2}} P_{i,\theta}\right\|^2\Big)\right) \le \log(\alpha).\end{aligned}$$

For each i, $t_i h_{i,\theta} + \frac{t_i^2}{2}\left\|\bar{\Sigma}^{\frac{1}{2}} P_{i,\theta}\right\|^2$ is a convex function of the optimization parameters θ. By [11, Example 3.14], given convex functions $g_1(\theta), \dots, g_r(\theta)$, the function $f(\theta) = \log\big(e^{g_1(\theta)} + \cdots + e^{g_k(\theta)}\big)$ is itself convex. It follows that $\log \mathbb{E}\big[\varphi(\eta_\theta)\big]$ is a convex function of θ and that the constraint set

$$\{\theta \in \Theta : \ \log \mathbb{E}[\varphi(\eta_\theta)] \le \log\alpha\} = \{\theta \in \Theta : \ \mathbb{E}[\varphi(\eta_\theta)] \le \alpha\}$$

is convex. Finally, from Lemma 9, if $\mathbb{E}[\varphi(\eta_\theta)] \le \alpha$ then $\mathbb{P}[\eta_\theta \le 0] \ge 1 - \mathbb{E}[\varphi(\eta_\theta)] \ge 1 - \alpha$. Together with Assumption 1, this implies that the optimization problem with constraint $\sum_{i=1}^r \exp\Big(t_i h_{i,\theta} + \frac{t_i^2}{2}||\bar{\Sigma}^{\frac{1}{2}} P_{i,\theta}||^2\Big) \le \alpha$ in place of (18) is a convex conservative approximation of (16)–(18). □

This convex approximation of Problem (16)–(18) is obtained in (13)–(15) by setting $R = 1$ and $\phi(\eta) = \varphi(\eta) - \alpha$. The result that the approximation is conservative relies essentially on the fact that, with this choice, $\phi(\eta) \ge \phi_{CC}(\eta)$ $\forall \eta \in \mathbb{R}^r$ (see Lemma 9). This result can be generalized: Given any two functions $\phi', \phi'' : \mathbb{R}^r \to \mathbb{R}$ such that $\phi'(\eta) \ge \phi''(\eta)$ $\forall \eta \in \mathbb{R}^r$, constraint (15) with $\phi = \phi'$ is more conservative than the same constraint with $\phi = \phi''$. This type of analysis can be exploited to compare different probabilistic constraints and to minimize the conservatism of the convex approximations with respect with the tunable parameters, but is not fully pursued here.

3.3 Ellipsoidal Constraint Functions

Consider the constraint function

$$\eta(\bar{x}_\theta, \bar{u}_\theta) = \left(\begin{bmatrix}\bar{x}_\theta \\ \bar{u}_\theta\end{bmatrix} - \delta\right)^T \Xi \left(\begin{bmatrix}\bar{x}_\theta \\ \bar{u}_\theta\end{bmatrix} - \delta\right) - 1,$$

where $\Xi \ge 0$ and δ are given. Then the constraint $\eta(\bar{x}_\theta, \bar{u}_\theta) \le 0$ restricts the vector $\begin{bmatrix}\bar{x}_\theta^T & \bar{u}_\theta^T\end{bmatrix}^T$ to an ellipsoid with center δ and shape determined by Ξ. We now provide an approximation of the chance constraint (18) that is a semi-definite program in the optimization parameters $\theta = (\bar{G}, \bar{d})$. Similar to §3.2.2, the idea is to ensure that the $100(1-\alpha)\%$ confidence ellipsoid of $\bar{w}$ is such that (18) holds. To this end, let

$$y_\theta = \begin{bmatrix}\bar{x}_\theta \\ \bar{u}_\theta\end{bmatrix} - \delta = \begin{bmatrix}\bar{A}x_0 + \bar{B}(\bar{G}\bar{w} + \bar{d}) + \bar{D}\bar{w} \\ \bar{G}\bar{w} + \bar{d}\end{bmatrix} - \delta = h'_\theta + P'_\theta \bar{w},$$

where $h'_\theta = \begin{bmatrix}\bar{A}x_0 + \bar{B}\bar{d} \\ \bar{d}\end{bmatrix} - \delta$ and $P'_\theta = \begin{bmatrix}\bar{B}\bar{G} + \bar{D} \\ \bar{G}\end{bmatrix}$.

Proposition 11 *Define $S_\theta = \beta(\alpha)\Xi^{\frac{1}{2}} P'_\theta \bar{\Sigma}^{\frac{1}{2}}$, with $\beta(\alpha)$ as in Proposition 7, and $\xi_\theta = \Xi^{\frac{1}{2}} h'_\theta$. Then*

$$\begin{bmatrix} -\lambda + 1 & 0 & \xi_\theta^T \\ 0 & \lambda I & (S_\theta)^T \\ \xi_\theta & S_\theta & I \end{bmatrix} \geq 0, \qquad \lambda > 0, \tag{26}$$

is a Linear Matrix Inequality (LMI) in the unknown parameters θ and λ. If (θ, λ) is a solution of (26), then θ satisfies (18). As a consequence, under Assumption 1, the problem

$$\begin{aligned} \inf_{\theta \in \Theta, \lambda} \quad & \mathbb{E}[V(\bar{x}_\theta, \bar{u}_\theta)] \\ \text{subject to} \quad & (7),\ (9) \text{ and } (26) \end{aligned}$$

is a convex conservative approximation of Problem (16)–(18).

Proof: The inequality (18) may be equivalently represented as

$$\begin{aligned} & \mathbb{P}\big(y_\theta^T \Xi y_\theta - 1 \leq 0\big) \geq 1 - \alpha \\ \iff \quad & \mathbb{P}\big(\|\Xi^{\frac{1}{2}} y_\theta\|^2 - 1 \leq 0\big) \geq 1 - \alpha \\ \iff \quad & \mathbb{P}\big(\|\Xi^{\frac{1}{2}} (h'_\theta + P'_\theta \bar{w})\|^2 - 1 \leq 0\big) \geq 1 - \alpha \\ \iff \quad & \mathbb{P}\big(\|\xi_\theta + S'_\theta \bar{w}\|^2 - 1 \leq 0\big) \geq 1 - \alpha, \end{aligned} \tag{27}$$

where $S'_\theta = \Xi^{\frac{1}{2}} P'_\theta$. Since $\bar{w} \sim \mathcal{N}(0, \bar{\Sigma})$, one can compute $\beta(\alpha) > 0$ such that $\left\|\bar{\Sigma}^{-1/2} \bar{w}\right\|^2 \leq \beta(\alpha)^2$ specifies the required $100(1-\alpha)\%$ confidence ellipsoid of $\bar{w}$. Hence, we need to ensure that $\left\|\bar{\Sigma}^{-1/2} \bar{w}\right\|^2 \leq \beta(\alpha)^2 \Rightarrow \|\xi_\theta + S'_\theta \bar{w}\|^2 \leq 1$. This is equivalent to

$$\sup_{\left\|\bar{\Sigma}^{-1/2} \bar{w}\right\|^2 \leq \beta(\alpha)^2} \|\xi_\theta + S'_\theta \bar{w}\|^2 \leq 1 \quad \iff \quad \sup_{\|\bar{v}\|^2 \leq 1} \|\xi_\theta + S_\theta \bar{v}\|^2 \leq 1.$$

It follows from [11, p. 653] that $\sup_{\|\bar{v}\| \leq 1} \|\xi_\theta + S_\theta \bar{v}\|^2 \leq 1$ if and only if there exists $\lambda \geq 0$ such that

$$\begin{bmatrix} -\xi_\theta^T \xi_\theta - \lambda + 1 & \xi_\theta^T S_\theta \\ (S_\theta)^T \xi_\theta & \lambda I - (S_\theta)^T S_\theta \end{bmatrix} \geq 0.$$

Using Schur complements the last relation can be rewritten equivalently as (26). Therefore, any solution of (26) implies (18). To verify that (26) is an LMI, note that S_θ and ξ_θ are affine in the optimization variables. Together with the assumed convexity of $\mathbb{E}[V(\bar{x}_\theta, \bar{u}_\theta)]$, the last statement of the proposition follows. □

4 Integrated chance constraints

In this section we focus on the problem

$$\inf_{\theta \in \Theta} \quad \mathbb{E}[V(\bar{x}_\theta, \bar{u}_\theta)] \tag{28}$$

$$\text{subject to} \quad (7),\ (9) \text{ and} \tag{29}$$

$$J_i(\theta) \leq \beta_i,\ i = 1, \ldots, r \tag{30}$$

where, for $i = 1, \ldots, r$, $J_i(\theta) \triangleq \mathbb{E}[\varphi_i \circ \eta_i(\bar{x}_\theta, \bar{u}_\theta)]$, functions $\varphi_i : \mathbb{R} \to \mathbb{R}$ are measurable and the $\beta_i > 0$ are fixed parameters. This problem corresponds to Problem (13)–(15) when setting $R = r$ and $\phi_i(\eta) = \varphi_i(\eta) - \beta_i$, with $i = 1, \ldots, r$. For the choice

$$\varphi_i(z) = \begin{cases} 0, & \text{if } z \leq 0, \\ z, & \text{otherwise,} \end{cases} \tag{31}$$

constraints of the form (30) are known as integrated chance constraints [19,20]. In fact, one may write (dropping the dependence of η_i on $\bar{x}_\theta$ and $\bar{u}_\theta$ to simplify the notation)

$$\mathbb{E}[\varphi_i(\eta_i)] = \mathbb{E}\left[\int_0^{+\infty} \mathbf{1}_{[0,\eta_i)}(s)\,\mathrm{d}s\right] = \int_0^{+\infty} \mathbb{P}\big(\eta_i > s\big)\,\mathrm{d}s,$$

where $\mathbf{1}_S(\cdot)$ is the indicator function of set S and the second equality follows from Tonelli's theorem [16, Theorem 4.4.5]. Therefore, constraint (30) is equivalent to

$$\int_0^{+\infty} \mathbb{P}\big(\eta_i(\bar{x}_\theta, \bar{u}_\theta) > s\big)\,\mathrm{d}s \leq \beta_i, \tag{32}$$

whence the name integrated chance constraint. Note that φ_i plays the role of a penalty (or barrier) function that penalizes violations of the inequality $\eta_i(\bar{x}, \bar{u}) \leq 0$, and β_i is a maximum allowable cost in the sense of (32). Of course, different choices of φ_i will not guarantee the equivalence between (30) and (32). However, they may be useful in deriving other quantitative chance constraint-type approximations.

4.1 Convexity of Integrated Chance Constraints

We now establish sufficient conditions on the η_i and φ_i for the convexity of the constraint set

$$\mathscr{F}_{\mathrm{ICC}} \triangleq \{\theta :\ J_i(\theta) \leq \beta_i,\ i = 1, \ldots, r\}. \tag{33}$$

The result is again a consequence of Proposition 3.

Proposition 12 *Let the mappings $\eta_i : \mathbb{R}^{(N+1)n \times Nm} \to \mathbb{R}$ be measurable and convex, and let the $\varphi_i : \mathbb{R} \to \mathbb{R}$ be measurable, monotone nondecreasing and convex. Assume that the $J_i(\theta)$ are finite for all θ. Then each $J_i(\theta)$ is a convex function of θ and $\mathscr{F}_{\mathrm{ICC}}$ is a convex set. As a consequence, under Assumption 1,* (28)–(30) *is a convex optimization problem.*

Proof: Fix $\omega \in \Omega$ arbitrarily. Since the mapping $\theta \mapsto \big(\bar{x}_\theta(\omega), \bar{u}_\theta(\omega)\big)$ is affine and $(\bar{x}, \bar{u}) \mapsto \eta_i(\bar{x}, \bar{u})$ is convex by assumption, their composition $\theta \mapsto \eta_i\big(\bar{x}(\omega), \bar{u}(\omega)\big)$ is a convex function of θ. Using the assumption that φ_i is monotone nondecreasing and convex, we may apply Proposition 3 with $\gamma(\omega, \theta) = \eta_i\big(\bar{x}_\theta(\omega), \bar{u}_\theta(\omega)\big)$ and φ_i in place of φ to conclude that $J_i(\theta) = \mathbb{E}[\varphi_i \circ \gamma(\omega, \theta)]$ is convex. Hence, for any choice of β_i, the set $\mathscr{F}_i \triangleq \{\theta :\ J_i(\theta) \leq \beta_i\}$ is convex. Since $\mathscr{F}_{\mathrm{ICC}} = \bigcap_{i=1}^{r} \mathscr{F}_i$, the convexity of $\mathscr{F}_{\mathrm{ICC}}$ follows. Together with Assumption 1, this proves that (28)–(30) is a convex optimization problem. □

It is worth noting that the function $\varphi(\cdot)$ of Section 3.2.4 satisfies analogous monotonicity and convexity assumptions with respect to each of the η_i, with $i = 1, \ldots, r$. Unlike those of Section 3, this convexity result is independent of the probability distribution of $\bar{w}$. By virtue of the alternative assumptions (iii′) and (iii″) of Proposition 3, the requirement that $J(\theta)$ be finite for all θ may be relaxed. A sufficient requirement is that there exist no two values θ and θ' such that $J(\theta) = +\infty$ and $J(\theta') = -\infty$. In particular, provided measurable and convex η_i, definition (31) satisfies all the requirements of Proposition 12.

Example 13 The (scalar) polytopic constraint function:

$$\eta(\bar{x}_\theta, \bar{u}_\theta) = T_x \bar{x}_\theta + T_u \bar{u}_\theta - y$$

fulfills the hypotheses of Proposition 12. Hence, the corresponding integrated chance constraint is convex.

Example 14 Following Example 1, an interesting case is that of ellipsoidal constraints. For an $\big((N+1)n+Nm\big)$-size positive-semidefinite real matrix Ξ and a vector $\delta \in \mathbb{R}^{(N+1)n+Nm}$, define

$$\eta(\bar{x}_\theta, \bar{u}_\theta) \triangleq \left(\begin{bmatrix}\bar{x}_\theta \\ \bar{u}_\theta\end{bmatrix} - \delta\right)^T \Xi \left(\begin{bmatrix}\bar{x}_\theta \\ \bar{u}_\theta\end{bmatrix} - \delta\right) - 1. \tag{34}$$

This is a convex function of the vector $\begin{bmatrix}\bar{x}_\theta^T & \bar{u}_\theta^T\end{bmatrix}^T$ (it is the composition of the convex mapping $\xi^T \Xi \xi$, $\Xi \geq 0$, with the affine mapping $\xi = \begin{bmatrix}\bar{x}_\theta^T & \bar{u}_\theta^T\end{bmatrix}^T - \delta$) and hence Proposition 12 applies.

Remark 15 A problem setting similar to Example 14 with quadratic expected-type cost function and ellipsoidal constraints has been adopted in [26], where hard constraints are relaxed to expected-type constraints of the form $\mathbb{E}[\eta(\bar{x}, \bar{u})] \leq \beta$. This formulation can be seen as a special case of integrated chance constraints with $\varphi(z) = z$ for all z. The choice of φ within a large class of functions is an extra degree of freedom provided by our framework that may be exploited to establish tight bounds on the probability of violating the original hard constraints, see Lemma 9 for an example.

4.2 Numerical solution of optimization problems with ICC

Even though ICC problems are convex in general, deriving efficient algorithms to solve them is still a major challenge [19,20]. For certain ICCs, it is possible however to derive explicit expressions for the gradients of the constraint function. Provided the cost has a simple (e.g. quadratic) form, this allows one to implement standard algorithms (e.g. interior point methods [11]) for the solution of the optimization problem.

Let $\bar{w}$ satisfy Assumption 2 with, for simplicity, $\bar{\Sigma}$ equal to the identity matrix, i.e. $\bar{w} \sim \mathcal{N}(0, I)$. Consider the problem with one scalar constraint (the generalization to multiple (joint) constraints is straightforward):

$$\begin{aligned} \min_{\theta} \quad & \begin{bmatrix}\bar{x}_\theta^T & \bar{u}_\theta^T\end{bmatrix} M \begin{bmatrix}\bar{x}_\theta \\ \bar{u}_\theta\end{bmatrix} \\ \text{subject to} \quad & (7),\ (9) \text{ and } \mathbb{E}[\varphi(T_x \bar{x}_\theta + T_u \bar{u}_\theta - y)] \leq \beta, \end{aligned} \tag{35}$$

where φ is defined as in (31).

Lemma 16 *Let z be a Gaussian random variable with mean μ and variance $\sigma^2 > 0$. Then*

$$\mathbb{E}[\varphi(z)] = \sigma g\Big(\frac{\mu}{\sigma}\Big)$$

where

$$g(x) = \frac{x}{2} \operatorname{erfc}\Big(\frac{-x}{\sqrt{2}}\Big) + \frac{1}{\sqrt{2\pi}} \exp\Big(\frac{-x^2}{2}\Big)$$

and $\operatorname{erfc}(x) = 1 - \operatorname{erf}(x)$ *is the standard complementary error function.*

Proof: Since $z \sim \mathcal{N}(\mu, \sigma^2)$ it holds that

$$\mathbb{E}[\varphi(z)] = \frac{1}{\sqrt{2\pi\sigma^2}} \int_0^\infty t \exp\left(-\frac{(t-\mu)^2}{2\sigma^2}\right)\, dt.$$

By the change of variable $y = \frac{t-\mu}{\sqrt{2}\sigma}$ one gets

$$\begin{aligned} \mathbb{E}[\varphi(z)] &= \frac{1}{\sqrt{\pi}} \int_{\frac{-\mu}{\sqrt{2}\sigma}}^\infty (\mu + \sqrt{2}\sigma y) \exp(-y^2) dy \\ &= \frac{\mu}{\sqrt{\pi}} \int_{\frac{-\mu}{\sqrt{2}\sigma}}^\infty \exp(-y^2) dy + \sigma\sqrt{\frac{2}{\pi}} \int_{\frac{-\mu}{\sqrt{2}\sigma}}^\infty y \exp(-y^2) dy \\ &= \frac{\mu}{2} \operatorname{erfc}\left(\frac{-\mu}{\sqrt{2}\sigma}\right) + \frac{\sigma}{\sqrt{2\pi}} \exp\left(-\frac{\mu^2}{2\sigma^2}\right), \end{aligned}$$

which is equal to $\sigma g(\frac{\mu}{\sigma})$. □

Simple calculations and the application of this lemma yield the following result.

Proposition 17 *Problem* (35) *is equivalent to*

$$\begin{array}{ll} \textit{minimize} & h_1(\bar{d}) + h_2(\bar{G}) \\ \textit{subject to} & \sigma g\Big(\dfrac{\mu}{\sigma}\Big) \leq \beta \end{array} \tag{36}$$

where

$$\begin{aligned} \mu =& (T^x \bar{A} x_0 - \bar{y}) + (T^x \bar{B} + T^u)\bar{d}, \\ \sigma =& ||(T^x \bar{D} + (T^x \bar{B} + T^u)\bar{G})^T||_2, \\ h_1(\bar{d}) =& \Big[(\bar{A}x_0 + \bar{B}\bar{d})^T \;\; \bar{d}^T\Big] \, M \begin{bmatrix} \bar{A}x_0 + \bar{B}\bar{d} \\ \bar{d} \end{bmatrix}, \\ h_2(\bar{G}) =& Tr\left(\Big[(\bar{B}\bar{G} + \bar{D})^T \;\; \bar{G}^T\Big] \, M \begin{bmatrix} \bar{B}\bar{G} + \bar{D} \\ \bar{G} \end{bmatrix} \right). \end{aligned}$$

Note that expectations have been integrated out. Now it is possible to put the problem in a standard form for numerical optimization. Let $\bar{G}_i$ and $\bar{D}_i$ be the i-th column of $\bar{G}$ and $\bar{D}$, respectively. Redefine the optimization variable θ as the vector $\bar{\theta} = \Big[\bar{d}^T \;\; \bar{G}_1^T \;\; \cdots \;\; \bar{G}_{Nn}^T\Big]^T$. Define e, f, v and L by

$$\begin{aligned} \mu = (T^x \bar{A} x_0 - \bar{y}) + (T^x \bar{B} + T^u)\bar{d} =& e + f^T \bar{\theta}, \\ (T^x \bar{D} + (T^x \bar{B} + T^u)\bar{G})^T =& v + L\bar{\theta}. \end{aligned}$$

Corollary 18 *Problem* (35) *is equivalent to*

$$\begin{array}{lll} \textit{minimize} & f_0(\bar{\theta}) = g_0(\bar{\theta}) + \displaystyle\sum_{i=1}^{Nn} g_i(\bar{\theta}) \\ \textit{subject to} & H\bar{\theta} = 0, \\ & f_1(\bar{\theta}) \leq 0 \end{array}$$

where the matrix H in the equality constraint accounts for the causal structure of $\bar{G}$, while

$$\begin{aligned} g_0(\bar{\theta}) =& \Big(\begin{bmatrix} \bar{A}x_0 \\ 0 \end{bmatrix} + \begin{bmatrix} \bar{B} \\ I \end{bmatrix} H_0\bar{\theta} \Big)^T M \Big(\begin{bmatrix} \bar{A}x_0 \\ 0 \end{bmatrix} + \begin{bmatrix} \bar{B} \\ I \end{bmatrix} H_0\bar{\theta} \Big), \\ g_i(\bar{\theta}) =& \Big(\begin{bmatrix} \bar{D}_i \\ 0 \end{bmatrix} + \begin{bmatrix} \bar{B} \\ I \end{bmatrix} H_i\bar{\theta} \Big)^T M \Big(\begin{bmatrix} \bar{D}_i \\ 0 \end{bmatrix} + \begin{bmatrix} \bar{B} \\ I \end{bmatrix} H_i\bar{\theta} \Big), \\ f_1(\bar{\theta}) =& \big\| v + L\bar{\theta} \big\| \, g\Big(\frac{e + f^T\bar{\theta}}{\big\| v + L\bar{\theta} \big\|} \Big) - \beta \leq 0, \end{aligned}$$

and H_0 and H_i are selection matrices such that $H_0\bar{\theta} = \bar{d}$ and $H_i\bar{\theta} = \bar{G}_i$.

We conclude the section by documenting the expressions of the gradient and the Hessian of the constraint function $f_1(\bar{\theta})$.

$$\nabla f_1(\bar{\theta}) = \frac{1}{\sqrt{2\pi}} L^T \frac{v + L\bar{\theta}}{\big\| v + L\bar{\theta} \big\|} \exp\Big(-\frac{(e + f^T\bar{\theta})^2}{2 \big\| v + L\bar{\theta} \big\|^2} \Big) + \frac{1}{2} \operatorname{erfc}\Big(-\frac{e + f^T\bar{\theta}}{\sqrt{2} \big\| v + L\bar{\theta} \big\|} \Big) f$$

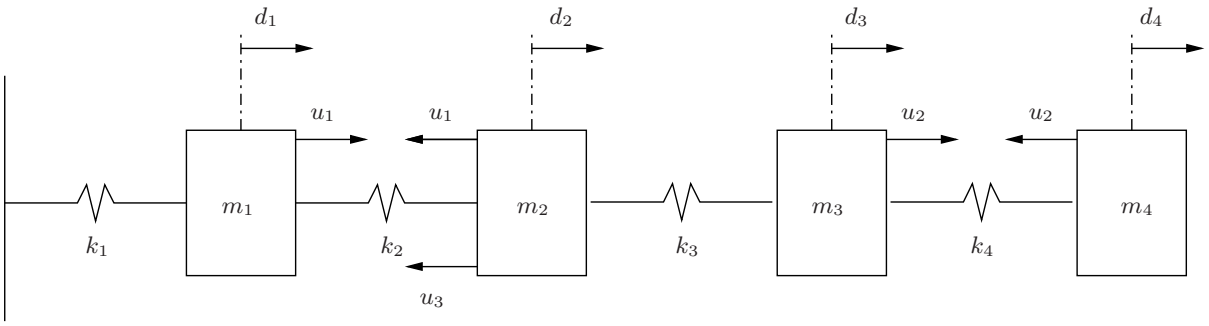

Fig. 2. A mechanical system with springs and masses.

$$\nabla^2 f_1(\bar{\theta}) = \frac{1}{\sqrt{2\pi}} \exp\Big(-\frac{(e+f^T\bar{\theta})^2}{2\left\|v+L\bar{\theta}\right\|^2}\Big)\Big[J_1(\bar{\theta}) + J_2(\bar{\theta}) - J_3(\bar{\theta})\Big]$$

where

$$\begin{aligned}
J_1(\bar{\theta}) &= \frac{1}{\left\|v+L\bar{\theta}\right\|}(L^TL + ff^T),\\
J_2(\bar{\theta}) &= \Big(\frac{(e+f^T\bar{\theta})^2 - \left\|v+L\bar{\theta}\right\|^2}{\left\|v+L\bar{\theta}\right\|^5}\Big)(L^T(v+L\bar{\theta})(v+L\bar{\theta})^TL),\\
J_3(\bar{\theta}) &= \Big(\frac{e+f^T\bar{\theta}}{\left\|v+L\bar{\theta}\right\|^3}\Big)(L^T(v+L\bar{\theta})f^T + f(v+L\bar{\theta})^TL).
\end{aligned}$$

The expressions of gradient and Hessian of the quadratic function $f_0(\bar{\theta})$, used e.g. by interior point method solvers, are quite standard and will not be reported here.

5 Simulation results

We illustrate some of our results with the help of a simple example. Consider the mechanical system shown in Figure 2. $d_1, \cdots, d_4$ are displacements from an equilibrium position, $u_1, \cdots, u_3$ are forces acting on the masses. In particular, u_1 is a tension between the first and the second mass, u_2 is a tension between the third and the fourth mass, and u_3 is a force between the wall (at left) and the second mass. We assume all mass and stiffness constants to be equal to unity, i.e. $m_1 = \cdots = m_4 = 1, k_1 = \cdots = k_4 = 1$. We consider a discrete-time model of this system with noise in the dynamics,

$$x(t+1) = Ax(t) + Bu(t) + w(t),$$

where w is an i.i.d. noise process, $w(t) \sim \mathcal{N}(0, \sigma^2 I), \sigma = 0.05$ for all t, and $x = \Big[d_1, d_2, d_3, d_4, \dot{d}_1, \dot{d}_2, \dot{d}_3, \dot{d}_4\Big]^T$. The discrete-time dynamics are obtained by uniform sampling of a continuous-time model at times $t \cdot h$, with sampling time $h = 1$ and $t = 0, 1, 2, \ldots$, under the assumption that the control action $u(t)$ is piecewise constant over the

sampling intervals $[t \cdot h, t \cdot h + h)$. Hence, $A = e^{hA_c}$ and $B = A_c^{-1}(e^{hA_c} - I)B_c$, where A_c and B_c, defined as

$$A_c = \begin{bmatrix} 0_{4\times4} & I_{4\times4} \\ A_{c,21} & 0_{4\times4} \end{bmatrix}, \qquad A_{c,21} = \begin{bmatrix} -2 & 1 & 0 & 0, \\ 1 & -2 & 1 & 0 \\ 0 & 1 & -2 & 1 \\ 0 & 0 & 1 & -1 \end{bmatrix},$$

$$B_c = \begin{bmatrix} 0_{4\times3} \\ B_{c,21} \end{bmatrix}, \qquad B_{c,21} = \begin{bmatrix} 1 & 0 & 0 \\ -1 & 0 & -1 \\ 0 & 1 & 0 \\ 0 & -1 & 0 \end{bmatrix}$$

are the state and input matrices of the standard ODE model of the system.

We are interested in computing the control policy that minimizes the cost function

$$\mathbb{E}\left[\sum_{t=0}^{N-1}\big(x(t)^T Q x(t) + u(t)^T R u(t)\big) + x(N)^T Q x(N)\right],$$

where the horizon length is fixed to $N = 5$, the weight matrices are defined as $Q = \begin{bmatrix} I & 0 \\ 0 & 0 \end{bmatrix}$ (penalizing displacements but not their derivatives) and $R = I$. The initial state is set to $x_0 = [0, 0, 0, 1, 0, 0, 0, 0]^T$. In the absence of constraints, this is a finite horizon LQG problem whose optimal solution is the linear time-varying feedback from the state

$$u(t) = -\big(B^T P(t+1) B + R\big)^{-1} B^T P(t+1) A x(t),$$

where the matrices $P(t)$ are computed by solving, for $t = N-1, \ldots, 0$, the backward dynamic programming recursion

$$P(t) = Q + A^T P(t+1) A - A^T P(t+1) B (B^T P(t+1) B + R)^{-1} B^T P(t+1) A,$$

with $P(N) = Q$. Simulated runs of the controlled system are shown in Figure 3. We shall now introduce constraints on the state and the control input and study the feasibility of the problem with the methods of Section 3. The convex approximations to the chance-constrained optimization problems are solved numerically in Matlab by the toolbox `CVX` [18]. In all cases we shall compute a 5-stage affine optimal control policy and apply it to repeated runs of the system. Based on this we will discuss the feasibility of the hard constrained problem and the probability of constraint violation.

5.1 Polytopic constraints

Let us impose bounds on the control inputs, $|u_1(t)| \leq 0.1$, $|u_2(t)| \leq 0.3$ and $|u_3(t)| \leq 0.15$, with $t = 0, \ldots, N-1$, and bounds on the mass displacements, $|d_i(t)| \leq 10$, for $i = 1, \cdots, 4$ and with $t = 1, \ldots, N$. In the notation of Section 3.2, these constraints are captured by the equation $\eta(\bar{x}, \bar{u}) = T^x \bar{x} + T^u \bar{u} - y \leq 0$ where $T^x = \begin{bmatrix} M^T & 0 \end{bmatrix}^T$ and $T^u = \begin{bmatrix} 0 & H^T \end{bmatrix}^T$, with

$$M = \begin{bmatrix} M_1 & & \\ & \ddots & \\ & & M_1 \end{bmatrix}, \quad M_1 = \begin{bmatrix} I_{4\times4} & 0_{4\times4} \\ -I_{4\times4} & 0_{4\times4} \end{bmatrix}, \quad H = \begin{bmatrix} I \\ -I \end{bmatrix},$$

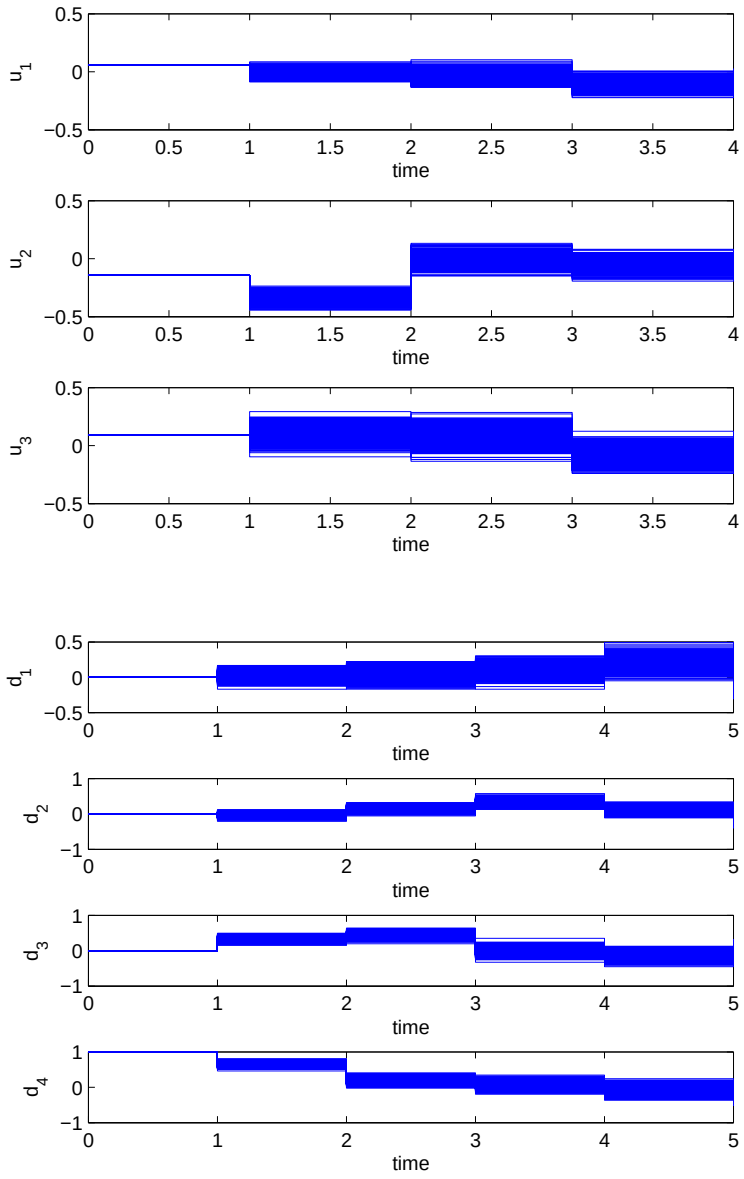

Fig. 3. 1000 sample paths of the system with LQG control policy. Above: control input. Below: mass displacements.

and

$$y = \begin{bmatrix} y_1 \\ y_2 \end{bmatrix}, \quad y_1 = \begin{bmatrix} 10 \\ \vdots \\ 10 \end{bmatrix}, \quad y_2 = \begin{bmatrix} y' \\ \vdots \\ y' \end{bmatrix}, \quad y' = \begin{bmatrix} 0.10 \\ 0.30 \\ 0.15 \end{bmatrix}.$$

This hard constraint is relaxed to the probabilistic constraint $\mathbb{P}[\eta(\bar{x}, \bar{u}) \leq 0] \geq 1 - \alpha$. The resulting optimal control problem is then addressed by constraint separation (Section 3.2.1) and ellipsoidal approximation (Section 3.2.2).

With constraint separation, the problem is feasible for $\alpha \geq 0.05$. For $\alpha = 0.1$, the application of the suboptimal control policy computed as in Proposition 6 yields the results shown in Figure 4. With this policy, the control input saturates within the required bounds whereas the mass displacements stay well below bounds. In fact, although the required probability of constraint satisfaction is 0.9, constraints were never violated in 1000 simulation runs. This

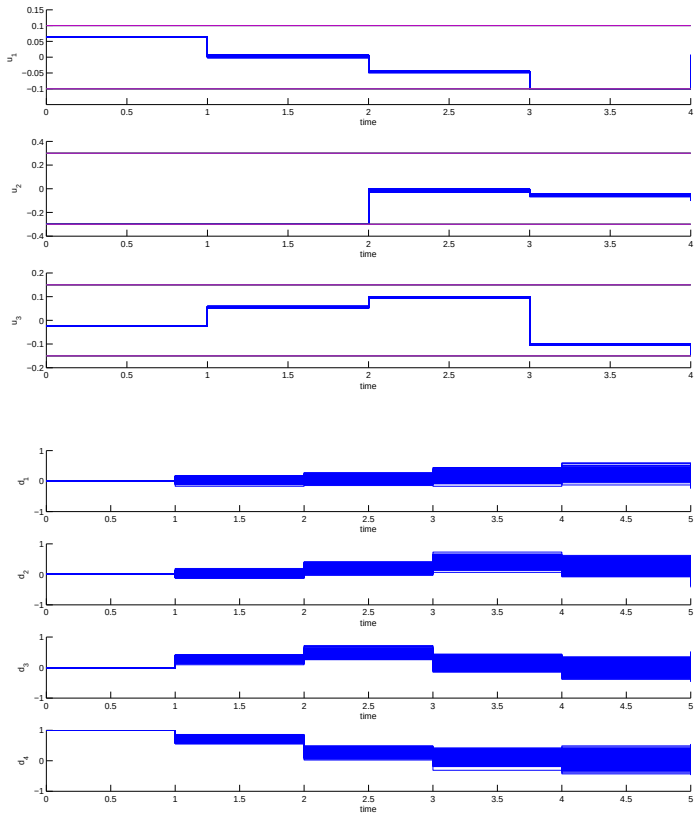

Fig. 4. 1000 sample paths of the system with control policy computed via constraint separation. Above: control input. Below: mass displacements. Horizontal straight lines show bounds.

suggests that the approximation incurred by constraint separation is quite conservative, mainly due to the relatively large number of constraints. It may also be noticed that the variability of the applied control input is rather small. This hints that the computed control policy is essentially open-loop, i.e. the linear feedback gain is small compared to the affine control term.

With the ellipsoidal approximation method, for the same probability level, the problem turns out to be infeasible, in accordance with the conclusions of Section 3.2.3. For the sake of investigation, we loosened the bounds on the mass displacements to $|d_i(t)| \leq 100$ for all i and t. The problem of Proposition 7 is then feasible and the results from simulation of the controlled system are reported in Figure 5. Although the controller has been computed under much looser bounds, the control performance is similar to the one obtained with constraint separation, a clear sign that the ellipsoidal approximation is overly conservative in this case. Another evidence of inaccuracy is the fact that, while the control inputs get closer to the bounds, the magnitude of the displacements is not reduced. As in the case of constraint separation, the applied control input is insensitive to the specific simulation run, i.e. the control policy is essentially open loop.

5.2 Ellipsoidal constraints

Consider the constraint function

$$\sum_{k=1}^{N} \Big(d_1(k)^2 + d_2(k)^2 + d_3(k)^2 + d_4(k)^2 \Big) + \|\bar{u}\|^2 \leq N \cdot c$$

with $c = (1^2 + 1^2 + 1^2 + 1^2 + 0.1^2 + 0.15^2 + 0.3^2) = 4.1225$. Unlike the previous section, we do not impose bounds on $d_i(k)$ and $u_i(k)$ at each k but instead require that the total "spending" on x and u does not exceed a total "budget".

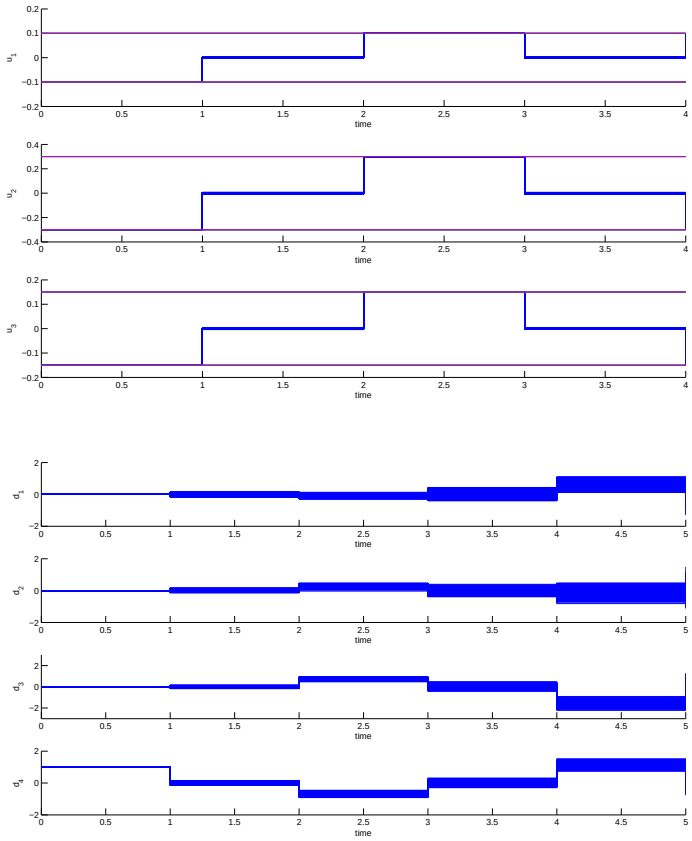

Fig. 5. 1000 sample paths of the system with control policy computed via ellipsoidal approximation. Above: control input. Below: mass displacements. Horizontal straight lines show bounds.

This constraint can be modelled in the form of Section 3.3, namely

$$\eta(\bar{x},\bar{u}) = \left(\begin{bmatrix}\bar{x}\\ \bar{u}\end{bmatrix} - \delta\right)^T \Xi \left(\begin{bmatrix}\bar{x}\\ \bar{u}\end{bmatrix} - \delta\right) - 1 \leq 0,$$

with $\delta = 0$ and $\Xi = \frac{1}{N\cdot c}\begin{bmatrix} L & 0 \\ 0 & I \end{bmatrix}$, where

$$L = \begin{bmatrix} L_1 & & \\ & \ddots & \\ & & L_1 \end{bmatrix}, L_1 = \begin{bmatrix} I_{4\times 4} & 0_{4\times 4} \\ 0_{4\times 4} & 0_{4\times 4} \end{bmatrix}.$$

The constrained control policy for $N = 5$ and $\alpha = 0.1$ is computed by solving the LMI problem of Proposition 11. Results from simulations of the closed-loop system are reported in Figure 6. Once again, constraints were not violated over 1000 simulated runs, showing the conservatism of the approximation. It is interesting to note that the displacements of the masses are generally smaller than those obtained by the controller computed under affine constraints, at the cost of a slightly more expensive control action. In contrast with the affine constraints case, the control action obtained here is much more sensitive to the noise in the dynamics, i.e. the feedback action is more pronounced.

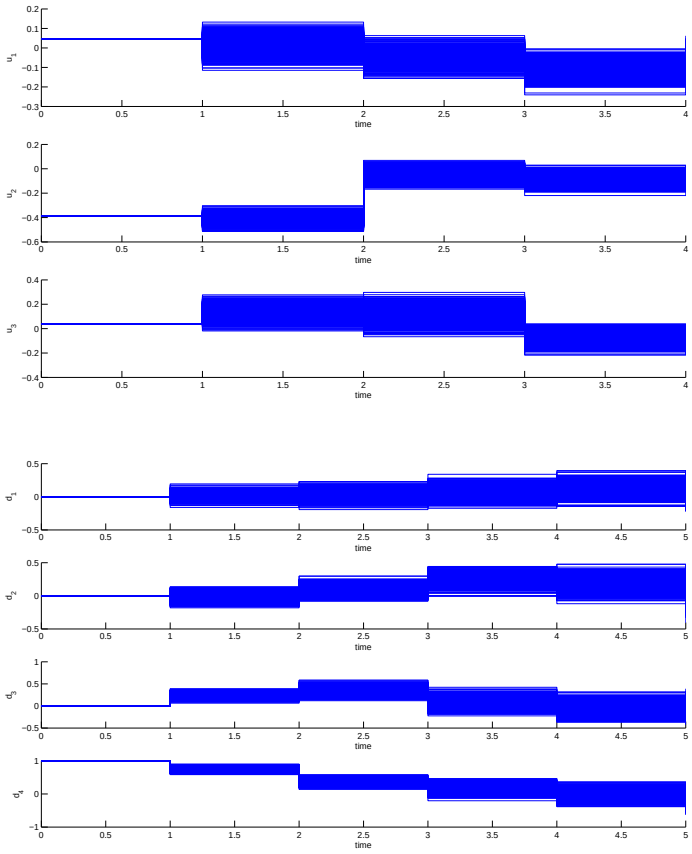

Fig. 6. 1000 sample paths of the system with control policy computed via ellipsoidal constraints. Above: control input. Below: mass displacements. Horizontal straight lines show bounds.

6 Conclusions

We have studied the convexity of optimization problems with probabilistic constraints arising in model predictive control of stochastic dynamical systems. We have given conditions for the convexity of expectation-type objective functions and constraints. Convex approximations have been derived for nonconvex probabilistic constraints. Results have been exemplified by a numerical simulation study.

Open issues that will be addressed in the future are the role of the tunable parameters (e.g. the α_i in Section 3.2.1, the t_i of Section 3.2.4 and the β_i in Section 4) in the various optimization problems, and the effect of different choices of the ICC functions φ_i (Section 4). Directions of future research also include the extension of the results presented here to the case of noisy state measurements, the exact or approximate solution of the stochastic optimization problems in terms of explicit control laws and the control of stochastic systems with probabilistic constraints on the state via bounded control laws.

References

[1] M. Agarwal, E. Cinquemani, D. Chatterjee, and J. Lygeros, *On convexity of stochastic optimization problems with constraints*, in Proceedings of the European Control Conference (ECC'09), 2009. Accepted.

[2] J. M. Alden and R. L. Smith, *Rolling horizon procedures in nonhomogeneous Markov decision processes*, Operations Research, 40 (1992), pp. S183–S194.

[3] K. Åström, *Introduction to Stochastic Control Theory*, Academic Press, New York, 1970.

[4] I. Batina, *Model predictive control for stochastic systems by randomized algorithms*, PhD thesis, Technische Universiteit Eindhoven, 2004.

[5] A. Ben-Tal, S. Boyd, and A. Nemirovski, *Extending scope of robust optimization*, Mathematical Programming, 107 (2006), pp. 63–89.

[6] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski, *Adjustable robust solutions of uncertain linear programs*, Mathematical Programming, 99 (2004), pp. 351–376.

[7] D. P. Bertsekas, *Dynamic programming and suboptimal control: A survey from ADP to MPC*, European Journal of Control, 11 (2005).

[8] D. Bertsimas and D. B. Brown, *Constrained stochastic LQC: a tractable approach*, IEEE Transactions on Automatic Control, 52 (2007), pp. 1826–1841.

[9] D. Bertsimas and M. Sim, *Tractable approximations to robust conic optimization problems*, Mathematical Programming, 107 (2006), pp. 5–36.

[10] V. I. Bogachev, *Gaussian Measures*, vol. 62 of Mathematical Surveys and Monographs, American Mathematical Society, Providence, RI, 1998.

[11] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, Cambridge, 2004.

[12] D. Chatterjee, E. Cinquemani, G. Chaloulos, and J. Lygeros, *Stochastic control up to a hitting time: optimality and rolling-horizon implementation*. `http://arxiv.org/abs/0806.3008`, jun 2008.

[13] D. Chatterjee, E. Cinquemani, and J. Lygeros, *Maximizing the probability of attaining a target prior to extinction*. `http://arxiv.org/abs/0904.4143`, apr 2009.

[14] D. Chatterjee, P. Hokayem, and J. Lygeros, *Stochastic model predictive control with bounded control inputs: a vector space approach*. `http://arxiv.org/abs/0903.5444`, mar 2009.

[15] P. D. Couchman, M. Cannon, and B. Kouvaritakis, *Stochastic MPC with inequality stability constraints*, Automatica J. IFAC, 42 (2006), pp. 2169–2174.

[16] R. M. Dudley, *Real Analysis and Probability*, vol. 74 of Cambridge Studies in Advanced Mathematics, Cambridge University Press, Cambridge, 2002. Revised reprint of the 1989 original.

[17] P. J. Goulart, E. C. Kerrigan, and J. M. Maciejowski, *Optimization over state feedback policies for robust control with constraints*, Automatica, 42 (2006), pp. 523–533.

[18] M. Grant and S. Boyd, *CVX: Matlab software for disciplined convex programming*, feb 2009. (web page and software) http://stanford.edu/∼boyd/cvx.

[19] W. K. K. Haneveld, *On integrated chance constraints*, in Stochastic programming (Gargnano), vol. 76 of Lecture Notes in Control and Inform. Sci., Springer, Berlin, 1983, pp. 194–209.

[20] W. K. K. Haneveld and M. H. van der Vlerk, *Integrated chance constraints: reduced forms and an algorithm*, Computational Management Science, 3 (2006), pp. 245–269.

[21] J. Löfberg, *Minimax Approaches to Robust Model Predictive Control*, PhD thesis, Linköping Universitet, Apr 2003.

[22] J. Maciejowski, *Predictive Control with Constraints*, Prentice Hall, 2002.

[23] A. Nemirovski and A. Shapiro, *Convex approximations of chance constrained programs*, SIAM Journal on Control and Optimization, 17 (2006), pp. 969–996.

[24] A. Prékopa, *Stochastic Programming*, vol. 324 of Mathematics and its Applications, Kluwer Academic Publishers Group, Dordrecht, 1995.

[25] J. Primbs, *A soft constraint approach to stochastic receding horizon control*, in Proceedings of the 46th IEEE Conference on Decision and Control, 2007, pp. 4797 – 4802.

[26] J. Primbs, *Stochastic receding horizon control of constrained linear systems with state and control multiplicative noise*, in Proceedings of the 26th American Control Conference, 2007.

[27] G. C. R. Tempo and F. Dabbene, *Randomized Algorithms for Analysis and Control of Uncertain Systems*, Springer-Verlag, 2005.

[28] J. Spall, *Introduction to Stochastic Search and Optimization: Estimation, Simulation, and Control*, Wiley, Hoboken, NJ, 2003.

[29] D. H. van Hessem and O. H. Bosgra, *A full solution to the constrained stochastic closed-loop mpc problem via state and innovations feedback and its receding horizon implementation*, in Proceedings of the 42nd IEEE Conference on Decision and Control, vol. 1, 9-12 Dec. 2003, pp. 929–934.

[30] M. Vidyasagar, *Randomized algorithms for robust controller synthesis using statistical learning theory*, Automatica, 37 (2001), pp. 1515–1528.